



Rec'd PCT/PFE 22 DEC 2004

ML 020626

Europäisches
Patentamt

European
Patent Office

Office européen
des brevets

#2
18/03/2529

10/519064

Bescheinigung

Certificate

Attestation

RECD 29 JUL 2003

WIPO PCT

Die angehefteten Unterlagen stimmen mit der ursprünglich eingereichten Fassung der auf dem nächsten Blatt bezeichneten europäischen Patentanmeldung überein.

The attached documents are exact copies of the European patent application described on the following page, as originally filed.

Les documents fixés à cette attestation sont conformes à la version initialement déposée de la demande de brevet européen spécifiée à la page suivante.

Patentanmeldung Nr. Patent application No. Demande de brevet n°

02077660.5

PRIORITY DOCUMENT
SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH
RULE 17.1(a) OR (b)

Der Präsident des Europäischen Patentamts;
Im Auftrag

For the President of the European Patent Office

Le Président de l'Office européen des brevets
p.o.

R C van Dijk

Best Available Copy



Anmeldung Nr:
Application no.: 02077660.5
Demande no:

Anmeldetag:
Date of filing: 02.07.02
Date de dépôt:

Anmelder/Applicant(s)/Demandeur(s):

Koninklijke Philips Electronics N.V.
Groenewoudseweg 1
5621 BA Eindhoven
PAYS-BAS

Bezeichnung der Erfindung/Title of the invention/Titre de l'invention:
(Falls die Bezeichnung der Erfindung nicht angegeben ist, siehe Beschreibung.
If no title is shown please refer to the description.
Si aucun titre n'est indiqué se referer à la description.)

Integer implementation of signal processing, in particular signal transforms

In Anspruch genommene Priorität(en) / Priority(ies) claimed /Priorité(s)
revendiquée(s)
Staat/Tag/Aktenzeichen/State/Date/File no./Pays/Date/Numéro de dépôt:

Internationale Patentklassifikation/International Patent Classification/
Classification internationale des brevets:

G06T1/00

Am Anmeldetag benannte Vertragstaaten/Contracting states designated at date of
filing/Etats contractants désignées lors du dépôt:

AT BE BG CH CY CZ DE DK EE ES FI FR GB GR IE IT LI LU MC NL PT SE SK TR

Integer implementation of signal processing, in particular signal transforms

BACKGROUND OF THE INVENTION

It is well known that common signal processing operations like FFTs, convolutions, DCT-transforms, etc., often require large dynamical ranges for the variables employed in such algorithms. This leads to implementations using *floating point* arithmetic, 5 rather than with *fixed point* arithmetic because the latter yield larger rounding noise at equal word-length. Let us first recall the distinction between integer or fixed-point arithmetic on the one hand and floating point arithmetic on the other.

Our goal is to represent numbers in the memory or registers of a computer or digital circuit in the form of binary digits ('0's and '1's). Because of their discrete nature we 10 can only represent a finite set of numbers, all other numbers are "rounded" or "truncated" to one of the representable values, leading to *quantization noise*. For the sake of the argument let us focus on numbers between -1.0 and 1.0, and say that we have 16-bits available to represent numbers in this range.

- for fixed point numbers, all representable numbers are of the form: $n \cdot 2^{-15}$, where n is an 15 integer in the range $[-2^{15} \dots 2^{15}-1]$. The representable numbers are uniformly spaced. The dynamic range, which is the ratio of the largest to the smallest representable number is $2^{15} \approx 10^5$.
- for floating point numbers, all representable numbers are of the form $s \cdot (0.5 + m/2^{a+1}) \times 2^{e-b}$, where m is an a -bit integer, and $(0.5 + m/2^a)$ is called the "mantissa" and obviously lies 20 between 0.5 and 1. s is a 1-bit "sign", and e is called the "exponent", a $(15-a)$ -bit number, and b is the exponent-bias. As an example take a 7-bit mantissa and an 8-bit exponent. Then the range 0.5...1 (set $e=b$) contains 128 representable numbers 1/256 apart. The range 0.25...1 (set $e=b-1$) also contains 128 representable number 1/512 apart, etc. We see that the representable numbers are packed closer together, the closer we get to 0, in a 25 logarithmic fashion. The exponent bias b sets the origin of this quasi-logarithmic scale. In this example the dynamic range is about $2^{256} \approx 10^{77}$.

Floating point numbers are a trade-off between a large dynamic range, and locally uniform distribution of representable numbers. This meshes nicely with the idea that in many relevant computations we need to represent small numbers with a small granularity

and large numbers with a large granularity. Another way to say this: the floating point representation matches the "natural distribution of numbers", which is roughly logarithmic, much more closely. For that reason, in practice floating point calculations almost invariably lead to much more accurate results than fixed point calculations with words of the same size (number of bits).

5 The major drawback of floating point numbers is that they require more complex hardware to perform additions, multiplication etc. E.g. for a floating point addition, both operands have to be normalized to the same exponent, followed by an ordinary addition, and a final re-scaling of the exponent. In software floating-point operations are therefore 10 usually much slower.

10 In the case of watermark detection, DSP operations like FFTs must happen accurately: a watermark is carefully hidden in the content -often in the LSBs- and so the signal processor must proceed with care so as not to lose it. However for watermarking in copy-protection or tracing applications, cost is a major issue: it is not a feature which can 15 warrant a higher price in the store. A manufacturer of watermark detectors has two choices to control the accuracy:

- use a floating point implementation with relatively high hardware cost (or high CPU-load for software)
- use a fixed point implementation, but considering the statements about accuracy above, 20 one is forced to use much longer words for an integer implementation than for a floating points. This also drives up the cost if many memory words are needed for storage, and consequently a lot of memory bandwidth is needed.

25 Applicant's International Patent Application WO 99/45707 discloses a watermark embedding system (hereinafter referred to as JAWS) to which the invention is particularly applicable. Fig. 1 shows the signal processing steps of this watermark detection. Since for algorithms like JAWS, memory was the largest cost-factor, up to now only the floating point implementation (with 17-point words: 8-bit mantissa, and an 8-bit exponent) is available. For instance, in the 2D FFT-step, if one wanted to use integers, one would need about 20...24 bits (depending on the video-content) to get similar accuracy to the 17-bit 30 floating-point implementation.

From the literature there are many methods known which can help to reduce the word-length for integer FFTs e.g.:

- insertion of guard-bits (shifting the input to the right by k bits, if the signal processing step is expected to increase the dynamic range by k -bits). An FFT on $N=2^n$ points

increases the dynamic range by N (worst case), so the required insertion of n guard bits is like pre-dividing the input by a factor of N .

5 - using block floating points. Block floating points are really (e.g. 16-bit) fixed points which represent a different range of numbers (like $[-1\dots 1]$, or $[-\frac{1}{4} \dots \frac{1}{4}]$ or $[-8\dots 8]$) at different stages in the processing step, depending on the required dynamic range. For instance in an FFT one would choose a new range for the 16-bit variables to represent, after every one of the n stages.

10 Although these methods are helpful, in general they still cause too much quantization noise to allow e.g. robust watermark detection.

10

OBJECT AND SUMMARY OF THE INVENTION

It is one of the purposes of this invention disclosure to show how it is possible to use fixed-point implementation to do signal processing with words which are no longer than the floating point implementation.

15

These and other objects are achieved by a method and arrangement as claimed in the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

20 Fig. 1 shows a block diagram indicating the signal processing steps involved to perform a watermark detection using the JAWS system.

Fig. 2 shows pictures illustrating the operation of the JAWS watermark detection system which is shown in Fig. 1.

Fig. 3 shows the power spectrum density of the output of the 2D FFT without (left) and with (right) pre-filtering in accordance with the invention.

25

Fig. 4 shows a diagram illustrating watermark detection reliability without and with pre-filtering in accordance with the invention.

Fig. 5 shows a block diagram of the JAWS watermark detection architecture including the pre-filter in accordance with the invention.

30

Fig. 6 shows a block diagram illustrating an algorithm for performing cyclical correlations using pre-filters to compress the dynamic range of the input and a post-filter to undo the effects of the pre-filter.

DESCRIPTION OF EMBODIMENTS

For many applications the statistics of the input-signal to the signal processing step is well known. For instance in the case of the watermark detector shown in Fig. 1, the output of the IDCT (input of the 2D FFT) contains essentially the original video-content which is strongly peaked in the low horizontal and vertical frequencies. This is illustrated in Fig. 2, where the left picture shows a typical input to the 2D FFT (here 1 second of folded video), and the right picture shows an intensity plot of the spectrum at the output of the 2D FFT, with a dynamic range of about 21 bits. The 0-freq. is at the center of the picture. The horizontal and vertical DC-frequencies cause the FFT to overflow. Preventing overflow by re-scaling yields even more quantization noise. This quantization noise ultimately obscures the watermark.

According to this invention it is preferred to first apply a (high-pass) filter to the input of the FFT which suppresses the low frequencies with large amplitude, thus reducing the required dynamic range. In fact the filter should be chosen such that it emphasizes those frequencies that contain most of the watermark energy, and cause least quantization noise in that energy range. As an example we have filtered the fold buffer (the left picture in Fig. 2) with the 2D filter:

$$\begin{bmatrix} 1 & -2 & 1 \\ -2 & 4 & -2 \\ 1 & -2 & 1 \end{bmatrix}$$

before inputting into the 2D FFT.

The result thereof is shown in Fig. 3. In this Figure, the left diagram shows the two-dimensional power density spectrum of the output of the 2D FFT without pre-filtering. The vertical axis is in powers of 2. It is clear that about 21 bits of dynamic range are required to represent the bulk of the coefficients. The right diagram shows the output of the 2D FFT with pre-filtering, using the filter mentioned above. The required dynamic range has been reduced to about 8-bits. The 2D spectra have been projected onto the plane given by $x+y=0$.

It is clear that the pre-filtering step causes a major decrease in required dynamic range to represent the (intermediate) result. In fact in combination with use of 16-bit block-floating point variables throughout the rest of the algorithm, the peak-reliability is almost indistinguishable from a 32-bit floating point implementation. Fig. 4 shows the watermark detection reliability for the same video sequence, processed once with a 32-bit floating point detector (solid line) and once with a 16-bit integer detector (dashed line).

Because of the normalization step (SPOMF's phase-extraction right after the 2D FFT), the effect of the input filter is purged. In other words: although with infinite precision variables the pre-filter has no effect (its effects are removed by SPOMF), it does reduce quantization noise significantly for a fixed-point implementation.

5

Fig. 5 shows a block diagram of the JAWS watermark detection architecture including the pre-filter in accordance with the invention to reduce dynamic-range / overflow / quantization noise problems in the 2D FFT. The advantages of pre-filtering for JAWS are:

10

- cheaper integer-only hardware / smaller CPU-load for integer-only software
- allows implementation on integer-only DSP (the vast bulk of DSP in use today, and certainly the only DSPs available in the low-end market; floating-point DSP are still very expensive).
- reduces memory consumption for JAWS detection by a factor 16/18 \times
- reduces memory bandwidth requirements significantly. Some implementors will not embed the 72 kbytes of RAM needed by the JAWS system on board of a chip but rather re-use 72Kbytes of an external DRAM memory chip to (like from a 512 kb drive-buffer in a DVD-ROM drive). Since the memory databus usually has a width of 16-bit, or 32 bits, an 18-bit transfer requires about (32/18) \times more data to move across the bus than necessary. This is particularly problematic in high-speed (16 \times) DVD-ROM drives, where memory bandwidth is at a premium.

15

Although the invention emerged out of research in the field of JAWS watermarking, it will be appreciated that the method is general enough to benefit other watermarking methods, or indeed signal processing in general. As an example, let us compute the (cyclical) correlation between the 2 patterns, *A* and *B*, by performing a multiplication in the frequency domain, as shown in Fig. 6. Again we may employ appropriately chosen pre-filters *F* and *G*, adapted to the statistical behavior of pattern *A* and *B* to control quantization noise in the FFTs. The effect of the filter (which scales the frequency components) is undone by the post-filter. If, as often is the case, pattern *B* is fixed, the post-filtering step may be combined with the pre-filtering and FFT of pattern *B*.

As another example, consider the fingerprinting technique described in the paper "Robust Audio Hashing for Content Identification," by Jaap Haitsma, Ton Kalker and Job Oostveen, presented at the Content-Based Multimedia Indexing conference 2001, Brescia, Italy. In this technique a "fingerprint" of an audio signal is generated by splitting its power spectrum in a series of frequency bands and coding differences between these bands, both with respect to frequency and with respect to time, in a small number of bits. The

fingerprint thus obtained is robust against a wide range of signal distortions, such as MP3 compression, noise addition, all-pass filtering, etc. Typically, the frequency bands considered cover the interval from 300 Hz to 3 kHz.

5 The power spectrum is obtained by applying a FFT to the downsampled and windowed input signal. As long as a floating-point algorithm is used this is just fine. However, quite often the power spectrum contains a peak near DC, which is substantially higher than the values in the frequency range of interest. This results in excessive quantization noise in that frequency range if an integer FFT with small dynamic range is used. Evidently, this may readily lead to spurious bit errors in the fingerprint, not caused by 10 actual signal distortions, but caused by a deficiency of the implementation. The solution is to remove the DC peak by applying a high-pass filter to the input signal prior to performing the FFT, or alternatively, to apply a band-pass filter which only selects the frequencies of interest.

15 The invention can also be described in the following manner. A digital signal processor operating with integer arithmetic circuits has a certain accuracy. Each processing step (multiplication, addition) increases the number of bits (the word length). For example, the Fast Fourier Transform having a butterfly structure requires a plurality of such processing steps to be performed. In practical implementations, the processing steps are recursively performed by a single integer arithmetic circuit having a given word length, say N. After 20 each step, the word length of the signal is reduced to the given word length N by rounding, truncation, or some other smart form of quantization. An obvious way to prevent quantization errors is to scale down the input signal. However, this results in quantization errors to be already introduced in the input signal. For processes such as watermark detection this is fatal, since the least significant bits of the input signal constitute precisely the place where the 25 watermark is embedded.

In accordance with the invention, the signal is pre-processed by a pre-processor which reduces the word length and which is invariant with respect to the subsequent process. The expression "being invariant" means that if the pre-processor operated with infinite accuracy, it would have no effect on the subsequent process. If such a 30 pre-processor operates with finite accuracy, it will reduce the quantization noise. The high-pass pre-filter described above with reference to the JAWS watermark detection process fulfills this condition, because it is a zero-phase filter and the watermark to be detected is carried by the phase of Fourier coefficients.

The invention can be summarized as follows. For some signal processing applications it is necessary to drive the cost down by choosing an integer-only implementation, with as small as possible bit-depth. However, often economical bit-depths like 16-bits yield inaccurate results because of excessive rounding noise. This holds in 5 particular for FFTs like they are used in watermark-detection algorithms like JAWS. For this reason such algorithms are forced to stick to 18-bit floating point arithmetic causing higher silicon-cost, memory bandwidth, and CPU load. To deal with this problem we propose to *pre-filter* the input signal prior to manipulating it, to control the dynamic range during the signal processing steps.

CLAIMS:

1. A method of processing a signal using a digital signal processor having a given word length, the method comprising the step of pre-processing said signal using a pre-processor which reduces the word length and performs an operation which is invariant with respect to the process being performed by the digital signal processor.

5

2. A method as claimed in claim 1, wherein said process being performed by the digital signal processor is watermark detection, and the pre-processor is a high-pass filter.

10

3. A method of processing a signal received in the form of signal samples having a range of sample values, the method comprising the steps of filtering the signal to reduce the range of signal sample values in a given band of non-interest, and digitally processing the filtered signal using integer arithmetic.

4. A digital signal processor comprising:

15

- input means for receiving a signal in the form of integer signal samples having a range of sample values;
- filtering means to reduce the range of signal sample values in a band of non-interest;
- a digital signal processing circuit for digitally processing the filtered signal using integer arithmetic.

20

5. A processor as claimed in claim 4, wherein said digital processing circuit comprises a transform circuit for transforming the signal into frequency coefficients.

25

6. A processor as claimed in claim 3, wherein said correlation circuit includes a Fourier transform circuit for computing said correlation for a plurality of shifts between said signal and said predetermined watermark pattern.

ABSTRACT:

For some signal processing applications it is necessary to drive the cost down by choosing an integer-only implementation, with as small as possible bit-depth. However, often economical bit-depths like 16-bits yield inaccurate results because of excessive rounding noise. This holds in particular for FFTs like they are used in watermark-detection 5 algorithms like JAWS. For this reason such algorithms are forced to stick to 18-bit floating point arithmetic causing higher silicon-cost, memory bandwidth, and CPU load. To deal with this problem we propose to *pre-filter* the input signal prior to manipulating it, to control the dynamic range during the signal processing steps.

10 Fig. 5

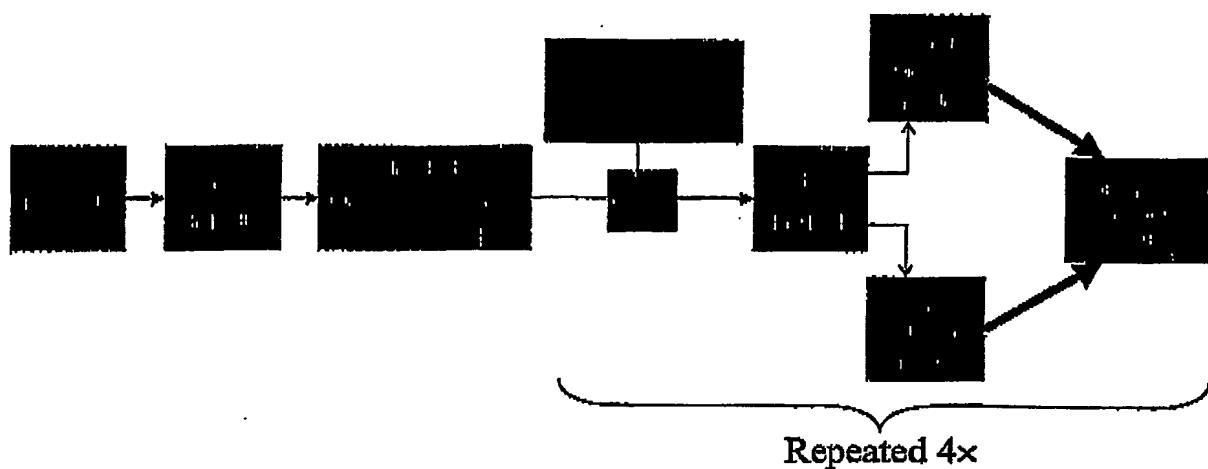


FIG. 1

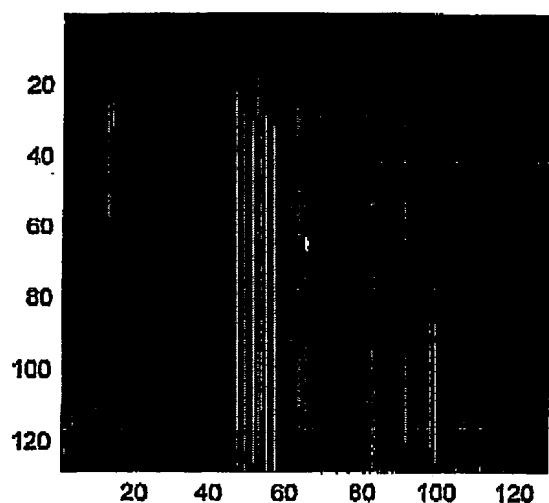
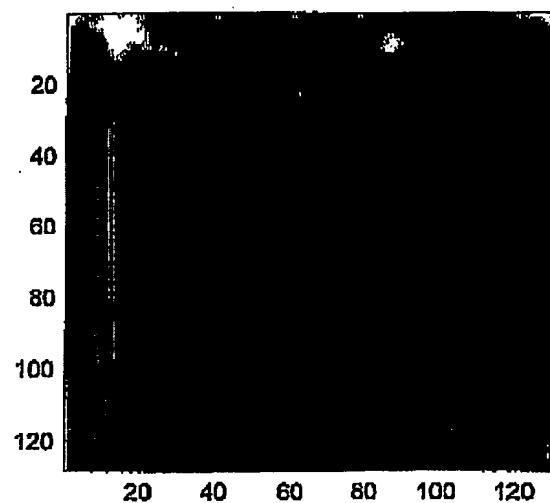


FIG. 2

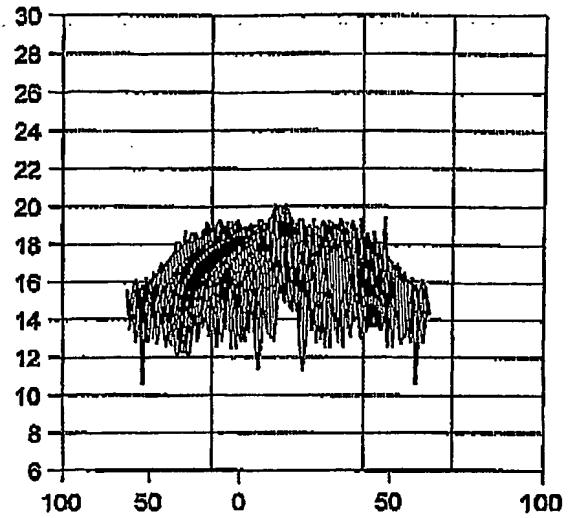
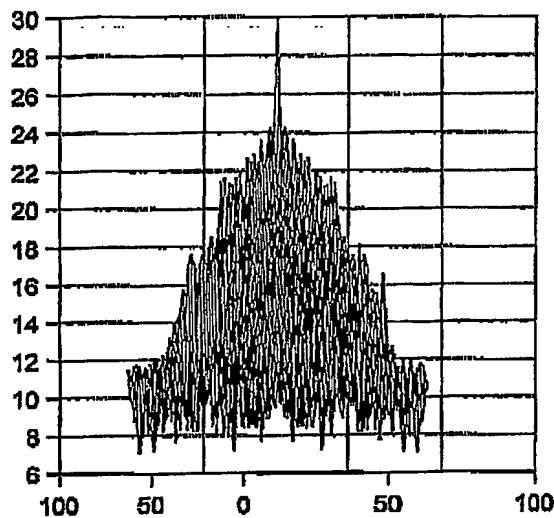


FIG. 3

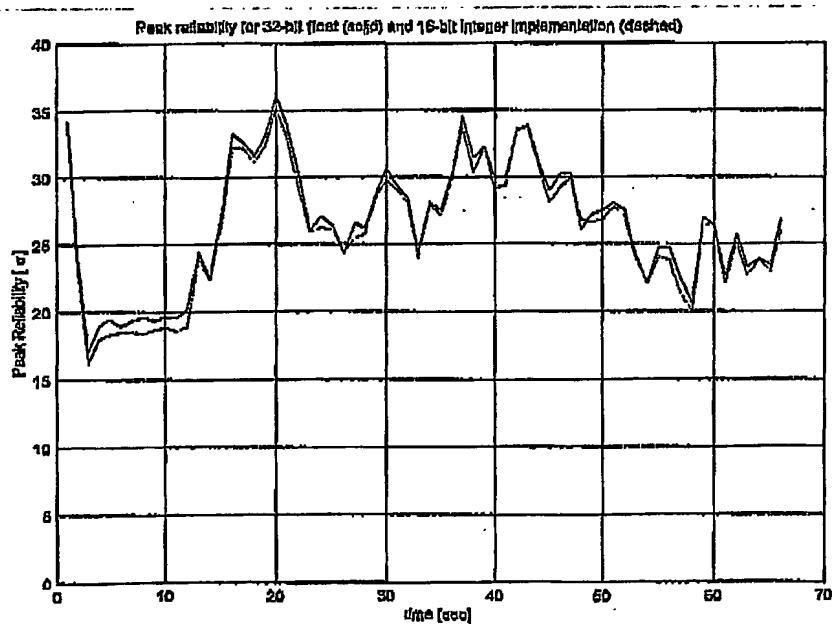
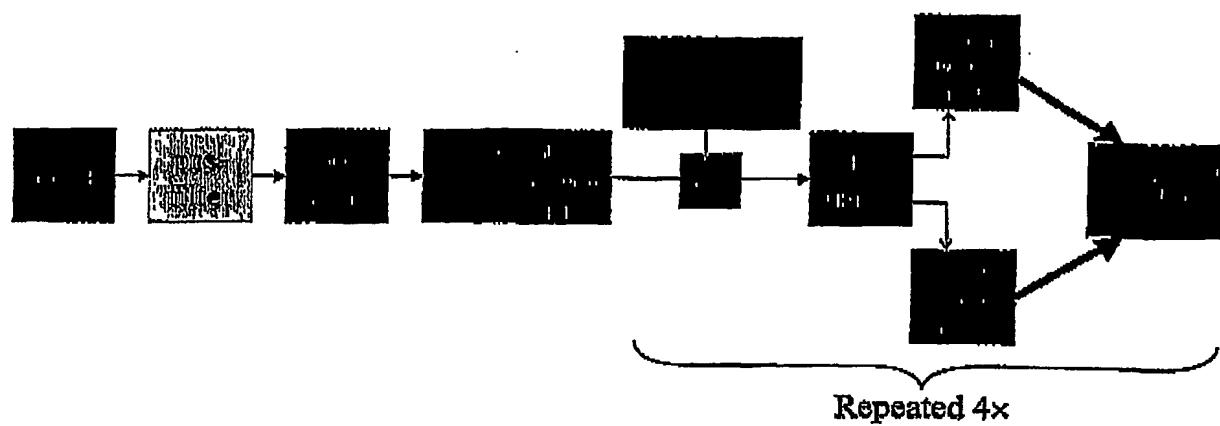
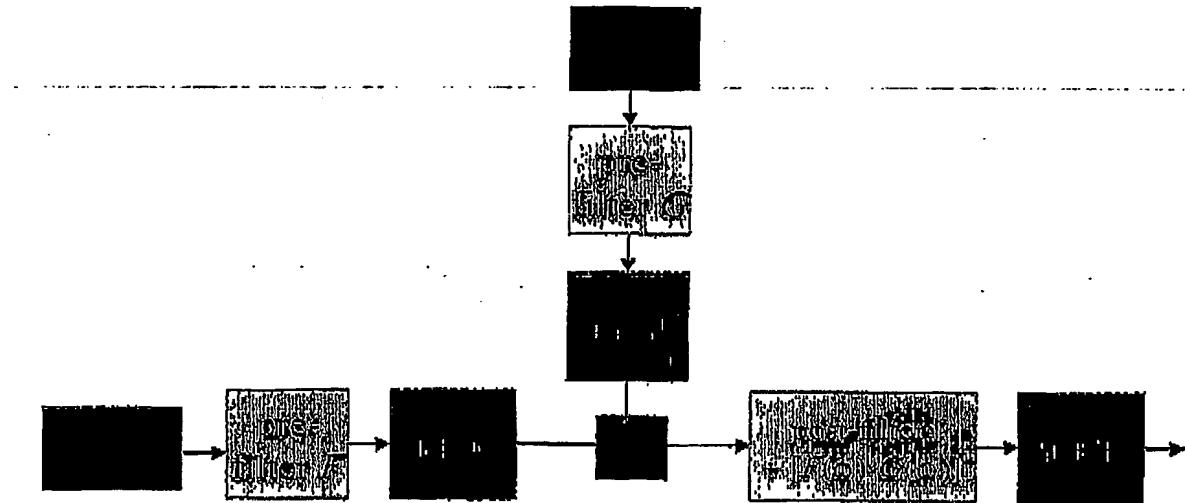


FIG. 4

**FIG. 5****FIG. 6**

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.